

#9 06-21-01



U.S. DEPARTMENT OF COMMERCE PATENT AND TRADEMARK OFFICE			
<b>DECLARATION UNDER 37 C.F.R. 1.132 OF DWIGHT A. MERRIMAN</b>		Docket Number: 11032/2144	
Reissue Applicant Dwight Allen MERRIMAN et al	Reissue Application No. 09/577,798	Reissue Filing Date May 24, 2000	
Patent Number 5,948,061	Issued September 7, 1999	Examiner	Art Unit 2757
Invention Title METHOD OF DELIVERY, TARGETING, AND MEASURING ADVERTISING OVER NETWORKS		Assignee DoubleClick Inc.	
<p>Address to: Assistant Commissioner for Patents Washington D.C. 20231</p> <p>I, Dwight A. Merriman, declare that:</p> <ol style="list-style-type: none"><li>1. I am presently Chief Technical Officer at DoubleClick Inc. ("DoubleClick"), assignee of the above identified patent and reissue application.</li><li>2. I am a co-inventor, together with Kevin J. O'Connor, of the claimed subject matter in the above identified patent and reissue application.</li><li>3. I am submitting this declaration to provide information on the development of our invention and other background about the relevant technology. The substance of this declaration includes overlap with my written testimony dated October 20, 2000 for the litigation, DoubleClick v. L90, 00 Civ. 2690 (DLC). That previous written testimony also addressed other subjects, such as the defendant L90's incorrect allegation of inequitable conduct. It is my understanding that the complete text of my previous written testimony is included in the materials to be provided to the Patent Office as part of these reissue proceedings.</li><li>4. I graduated from Miami University in Ohio in 1989 with a Bachelor of Science Degree in Systems Analysis. After graduation, I worked as a software developer for a small company called Solid Software. At this time I was proficient in several computer languages, including Modular 2, C, C++ and Pascal.</li></ol>			

### **Background on the Relevant Technology**

5. Web pages are generally constructed in a language called HyperText Markup Language, or HTML. HTML is essentially a programming language designed to instruct Web browsers how to display a Web page for a user. When a user accesses a Web site, the user's browser makes a request to the site for the desired Web page. The page contains HTML, which is interpreted by the browser, with the result that the page is displayed as intended. The user does not see the HTML, but rather the collection of text and graphics elements that the browser was told to construct based on the HTML it received.
6. The basic premise of our technology (which we call DART, for Dynamic Advertising, Reporting and Targeting) is that the Internet enables Web sites to use elements from across the Internet to build pages for a user to see when he or she visits a particular site.
7. This is easiest to see by way of example. Consider a Web site set up by members of a family. The Web site is set up by the parents, who live in New York. They provide the system that "hosts" the site, and the site is accessed by connecting to their host computer, called a server. These parents have two children, who go to college in San Francisco and Chicago, and who each have their own Web sites hosted on their colleges' servers. The parents would like their Web site to act like on-line family albums, and display pictures of their children in college.
8. There are two ways to accomplish this. First, the children could send the pictures to their parents to post on their site, not unlike posting the pictures onto a bulletin board. In that case, the data for the pictures resides on the parents' computer, and, when a visitor accesses the parents' Web page containing the picture, the HTML instructs the browser to load the picture from the parents' system. Technically speaking, the Web page relays an instruction to the user in the form of an HTML `<img src>` (image source) tag, which tells the user where on the parents' computer the picture can be found.
9. However, rather than host the pictures themselves, the pictures could be located on the children's Web sites at college. The structure of the parents' Web page is the same, except now the `<img src>` tag tells the user to get the picture from the college Web server. The page displayed for the user is the same in either case.
10. This ability to assemble elements from different places is an inherent feature of the Web. What Kevin and I learned, however, is that we could use the `<img src>` feature of the Web to target advertising. The "image" retrieved by the tag did not have to be preselected – we could, and did, build an advertising system that treats a request based on an `<img src>` tag as an instruction to select an appropriate, targeted image, and deliver that image to the user.

11. Returning to the example above, assume that the parents would like to show advertising on their Web site, and wanted to use DoubleClick to provide those ads. Just as with the pictures of their children, the parents include in their site an `<img src>` tag for the advertisement, which refers to DoubleClick's ad server. However, rather than provide a pre-selected picture, like the requests to the college servers for the family pictures, the DoubleClick server analyzes the request and chooses an ad based on appropriate criteria. For example, if the parents' site was visited by friends of the child in San Francisco, DoubleClick might provide an ad tailored to San Francisco, whereas if it were accessed by friends of the child in Chicago, it might provide an ad tailored to that location.
12. By separating control of advertisement from the place where it is to be shown, many things became possible. First, selling ads across multiple sites is possible, by aggregating the sites containing advertisers into a network. So, if dozens of parents across the country set up Web pages like those in our example, DoubleClick could collect a body of ads and target them appropriately to visitors to those sites. The sites themselves would not necessarily know the ad that would get shown to visitors to their sites – participation in this network means that DoubleClick's system decides which ads should be shown and provides those ads directly to the user.
13. The advantage to advertisers in such a system is clear – they can provide advertising to a broad range of sites while dealing just with DoubleClick. Similarly, Web sites can receive ads from a broad range of advertisers (and derive revenue from those advertisements) with only one contact. Moreover, because they are part of a larger network, they have the opportunity to attract more advertising than they ever would alone – the system effectively increases their marketability. The network idea is key to our technology and our business model.
14. The advantages of DoubleClick's system were countered, in our initial development, by a difficult technical problem. In ordinary Web advertising, the Web page provides both the advertisement and the instructions to the user's browser telling it to link the user's browser to the advertiser's own Web site when the user "clicks" on the banner. But when the advertisement is provided from a third-party like DoubleClick, the Web page does not know who the advertiser is, and so cannot provide the link to the advertiser's page. To solve this problem, we devised the system so that, when we deliver an ad, we record identifying information for the user's browser (his IP address or "cookie" such that the user remains anonymous) and identifying information for the advertiser. So, when the user clicks on the banner ad, we can identify the user's browser and know that he or she is asking to be taken to the advertiser's site, and we then provide the Web address of the advertiser to the user. This system not only enables the browser to reach the advertiser, since the request passes through DoubleClick's ad server, but also it enables DoubleClick to track the numbers of anonymous users who click on its banners.

15. Beneath these basic components is a complex architecture that actually enables our system to perform these functions. With each request for an ad, we receive a great deal of anonymous information about the computer making the request. For example, we receive information telling us the computer's address on the network, the operating system it is using, and other technical information. This information is provided as part of the standard operation of the Internet.
16. Additional anonymous information is learned from the computer's address on the network (that is, its IP address), which will assist in the targeting of ads. Part of the address is unique to each user, but part of the address is used to identify the user's "network" or "domain" – that is, the service through which it is accessing the Internet. For example, people from DoubleClick accessing the Internet have, as part of their IP address, information that identifies them as accessing the Internet through *doubleclick.net*. Persons accessing through a Federal Court's system could be identified as accessing the Internet through *www.uscourts.gov*.
17. The IP address itself is a string of numbers that convey no particular information on their own. We designed the DoubleClick system to call upon various public databases to determine the domain name (e.g. *www.uscourts.gov*) relating to the network IP addresses we receive. Using other public databases, we can determine the geographic locations of these networks, and, for networks that appear to be businesses, we can learn the Standard Industry Code that will tell the system what general industry the business is in. This information is retrieved and compiled into a proprietary DoubleClick database based on the network IP address.
18. The DoubleClick system also builds a database containing records relating to each individual Web browser accessing the system. We do this by using the ability of the popular Web browsers (Netscape Navigator and Internet Explorer) to place "cookies" on a user's computer. A "cookie" is little more than an identifier that can be recognized by a specific Web site when it is accessed by a particular Web browser.
19. Generally, the first step in employing cookies occurs when the user places a request to a Web site. That Web site responds with the requested Web page, as discussed above, and also an instruction to the Web browser to give the user a cookie – that is, to provide the user's computer with a specific identifying code, known only to the user's computer and the Web site, that allows the Web site to recognize the user's computer uniquely.
20. When the user's browser next goes to a Web site from which it previously received a cookie, the browser transmits the cookie for that site to the Web site's servers. The servers then can identify the browser. This is used, for example, by sites to provide customized pages and service to a user, and to track a user's use of the site. Cookies are provided only to the site that originally placed them – thus, the New York Times Web site cannot access the user's Wall Street Journal cookie.

21. Our implementation of cookies, however, takes this aspect of Web browsing to a new level. As noted above, DoubleClick receives a request for an ad from the user's browser whenever a user visits a site in the DoubleClick Network of Web sites. Each time an ad request is made by a particular user, that DoubleClick cookie stored by that browser is provided to DoubleClick, regardless of which site in the network the user is visiting. DoubleClick then has two pieces of information from which to identify the user – the IP address, which is looked up in DoubleClick's network IP address database as discussed above, and the user's anonymous unique cookie. These two features allow DoubleClick to bring its full ability to target ads to bear at whatever site the user visits in the network of sites having their ads served through DoubleClick.
22. One result of this process is our ability to control the number of times a user sees a particular advertisement. Studies have shown that the attractiveness of an advertisement to a user decreases with the number of times that user has seen a particular ad. Since our advertisers pay by the impression – the number of times an ad is shown overall – they have an interest in limiting the number of times the ad has been shown to each individual user. By keeping track of which ads are shown to each user, and recognizing users' browsers by their cookies, we can ensure that each user does not see a particular ad more than, say, five times, regardless of where in the DoubleClick Network of sites they might have seen those ads.
23. This ability to target different ads across multiple Web sites is the core of the functionality that Kevin and I invented. We developed components capable of performing these functions quickly, efficiently, and reliably. The scope of the process is significant: today, the system receives 36,000 requests for ads per second (2.5 billion per day), and, for each request, must select the best ad to show from among 470,000 active ad placements. And, our system must complete this selection process in 20 milliseconds – while simultaneously handling other requests for ads and the click through process discussed above. Although initially we were able to provide targeted ad serving and click throughs by using one server connected to the Internet, today, roomfuls of interconnected computers are necessary to provide this functionality as demanded by the popularity and success of the service.
24. Finally, our system also is integrated with management, reporting, and billing functions. This system enables our nearly 5000 advertisers to work with us to set up ad campaigns and review the success of their advertising. Our over 3200 affiliated Web sites can review the impressions served on their sites (and thus the revenue they will receive). The system also enables us to perform billing for the ads we serve.
25. The foundations of the system set forth in the preceding section were developed by myself and Kevin in the fall of 1995, and are the basis for the disclosure and claims of the '061 Patent. In the following sections, I discuss how these ideas came into being.

### Creation of Internet Advertising Network

26. In 1990, I began working for Intercomputer Communications Corporation ("ICC"), a telecommunications software company that was acquired by Digital Communications Associates ("DCA") in Atlanta in 1992. DCA was then acquired by Attachmate in late 1994. I started as a software developer at ICC/DCA, but subsequently moved to the research and development team as a senior research engineer developing OpenMind, a computer group work product. I also led the design and development of Attachmate's first Internet applications, including a Web browser and HTTP server software.
27. I met Kevin O'Connor shortly after I started working at ICC. Kevin was one of the co-founder's of ICC and subsequently in charge of DCA's research and development team after DCA's acquisition of ICC. Kevin and I worked together on the development of the OpenMind product, and I reported directly to him until his departure from Attachmate in January 1995.
28. During the summer of 1995, Kevin and I began to discuss different ideas for starting our own company. In August, I decided to leave Attachmate to pursue a venture with Kevin, who had had substantial success in his own technology ventures previously. After I left Attachmate, we worked all day, seven days a week, in the basement of Kevin's house in Atlanta thinking about different ideas for a start-up venture. We spent a lot of time trying to figure out trends which could be overlaid with customer needs and technologies to form a product. One of these trends was Internet advertising, and, by early September, we came up with the idea that would eventually evolve into Internet Advertising Network.
29. We initially called our business the "Internet Advertising Federation" ("IAF"). However, in the Fall of 1995, we changed the name to Internet Advertising Network ("IAN"). For purposes of this declaration, I will refer to our business as IAN from the period of formation of IAF through our merger with the DoubleClick Division of Poppe Tyson (the "DoubleClick Division") in January 1996. This merger resulted in the creation of DoubleClick Inc. After we changed the name of the Internet Advertising Federation to IAN, we continued to use the IAF acronym, but applied it to a corollary service: the Internet Address Finder, which was an Internet directory of e-mail addresses. Any subsequent references to IAF in this declaration are to this corollary service.

30. At the beginning of September 1995, the Internet was still relatively nascent, and we spent a lot of time thinking about whether the economic model that was going to govern content distribution on the Internet would be based on a subscription or advertising model. At first we were skeptical of advertising and thought that the model would be subscription based. As such, we started thinking about a business that would aggregate content from multiple publishers into a network of Web sites so that users would not have to subscribe to multiple Web sites to get content from the Internet. However, as we began to think more about it, we became convinced that advertising would be the better model because, with advertising, content would be free for Internet users. As we thought more about advertising on the Internet, we realized that our initial concept of aggregating content from multiple publishers into a network of Web sites would be the right model for Internet advertising. By aggregating many Web sites into a network, we realized that we would be able to present advertisers with an larger overall audience and thereby significantly increase the efficiency of Internet advertising.

#### **Development of IAN's Technology**

31. After we conceived of the idea of aggregating a network of Web sites to which we would serve ads as a third party, I spent the month of September 1995, researching the theoretical feasibility of designing a system capable of executing third-party targeted advertising over the Internet. Through my research and discussions with Kevin about technology, Kevin and I conceived of the idea of targeting ads to individual users based on information gleaned from that user's Internet browser. By developing software that could extract information about a user accessing a network Web site through the use of Internet cookies and IP addresses and by simultaneously building a comprehensive database of user and domain demographic information, we could significantly increase the efficiency of Internet advertising by serving ads targeted to individual Internet users.
32. My first task was to figure out what sort of technology we would need to be able to serve ads as a third party to a network of Web sites. I had to determine how the system would be constructed and what its features and functionality would be. I remember facing serious theoretical and practical challenges to the development of the technology at that time. For example, I was uncertain whether I would be able to design a system that could respond to a user's click-through to an anchor tag on a Web site for which we would be trying to serve ads for multiple advertisers in a single ad spot. It was a significant problem to figure out whether we would be able to design a system that could send a user to a particular advertiser's Web site after such a click. After doing a lot of research and looking at a lot of different Web protocols and the recent implementation of cookies in popular Web browsers, I was able to figure out how to solve this particular problem.

33. By the end of September 1995, Kevin O'Connor and I had conceptually mapped out the majority of the core concepts of our invention. The basic theoretical construct of IAN's system architecture was set forth in our first business plan, which was last modified on September 25, 1995 and prepared in the ordinary course of business. (Attached as Exhibit A.) However, this construct only set out what the pieces of the system would be. It would take at least another two months to fully design and implement a prototype capable of serving targeted third-party ads over the Internet.
34. In October 1995, I really began the process of building the different pieces of the ad management system and starting to put it together, one piece at a time. Nobody had ever created an Internet system like the one we had designed, and it was extremely challenging for me to try and build the system by myself. Even working approximately 100 hours a week, the process was slow and arduous.
35. I recall doing some work in October and November 1995 building the derive profile process. This was the piece that would allow IAN to extract information about users, which would then be compiled in our database and allow us to serve targeted advertising. For example, at that time we were still working on logic flows and specifications for the targeting system. (See Exhibit B, which is a memo discussing some of our plans for logic flows and targeting ideas, created in the ordinary course of business.) Exhibit B represents an intermediate step in our development from conception to coding. As that document shows, we were still working out a way to "allow for multiple ads per page" and to establish unique identifiers for ads and requests for ads.
36. I wrote the source code for the ad server throughout the Fall of 1995 and then a lot of very significant work on the ad server source code was done in December and January. (A true and correct copy of one of the earliest existing versions of the ad server source code, created in the ordinary course of business, has been attached as Exhibit C.) That document, printed in 1998 in connection with the prosecution of the '061 Patent, bears a few dates beginning in September 1995, but the bulk of the code is dated in December 1995 and January 1996.
37. The ad server is the component of the ad management system that receives ad requests and makes real-time selections of ads and determines which ad for multiple advertisers should be delivered. It was not until December that I had written the minimal amount of code necessary for the ad server to serve a targeted third-party ad over the Internet. This memory is supported by a document outlining the status of various IAN projects that was last modified on December 7, 1995. (A true and correct copy of this document, which was created in the ordinary course of business, has been attached as Exhibit D.) This exhibit, which represents a snapshot of the state of IAN on December 7, 1995, demonstrates that as of that date, the ad server was only 80% complete.

38. My recollection is further supported by several e-mails that I sent to Kevin O'Connor in December 1995. For example, on December 12, 1995, I wrote: "Making good progress on the ad server, making it more robust. Ready for targeted ads." (A true and correct copy of this e-mail, which was prepared in the ordinary course of business, has been attached as Exhibit E.) The next day, I confirmed by e-mail that SIC targeting, based on Standard Industry Codes, was functional. (A true and correct copy of this e-mail, which was prepared in the ordinary course of business, has been attached as Exhibit F.)
39. We served our first targeted test ads sometime after the major work had been completed on the ad server in December. My recollection is supported by a contemporaneous e-mail dated December 21, 1995 which states: "There are now four ads running on the production ad server: . . . 1. ISS, untargeted . . . 2. ISS, targeted to Unix . . . 3. IAN jobs, targeted to Georgia . . . 4. Tell a friend about IAF, untargeted." (A true and correct copy of this e-mail, which was prepared in the ordinary course of business, has been attached as Exhibit G.)
40. This e-mail indicates that as of December 21, 1995, IAN started serving ads for Internet Security Systems ("ISS") that were targeted to users using Unix based systems and ads for IAN targeted to Internet users from Georgia. We were also serving "untargeted" ads for ISS and the Internet Address Finder. "Untargeted" ads were ads that were hard-coded into the IAF Web site and which were presented to users without consulting any characteristics of the individual users.
41. By the beginning of December, Scott Leonard, another former Attachmate employee, joined as an additional engineer at IAN. From December through March, Scott and I continued to work around the clock writing code and strengthening the system to prepare for our March 1 goal of launching a beta-test version of the network.
42. By the middle of January 1996, I believe most forms of ad serving and targeting supported by the early system were ready for testing. I have reviewed an e-mail I sent, in the regular course of business, to Thomas Wharton of Poppe Tyson and Kevin O'Connor, where I describe our plan for functionality testing and stress testing of the system. (A true and correct copy of this e-mail, which was prepared in the ordinary course of business, has been attached as Exhibit H.) As part of the test, I planned to test various ads with different targeting criteria and frequency control. This leads me to believe that we were still testing our functionality as of this point, which we needed to do "live" on the Internet. Moreover, it reminds me that we had contemplated a stress test, to determine the ability of our ad server to satisfy high demand, but that as of this point in time we did not even have fast enough Internet connectivity to support that kind of a test.
43. After the January 1996 merger between IAN and the DoubleClick Division of Poppe Tyson (see above), DoubleClick's sales force began to solicit Web sites to be a part of DoubleClick's beta-test of the third-party ad serving network.

44. The DoubleClick beta-test began on March 1, 1996. When the beta-test began, we had about 10 Web sites that had signed on to be part of the test of the network. While the test went well, it demonstrated that we still had some bugs to work out of the system before we would be able to go live. We spent the month of March perfecting the system and officially launched the DoubleClick network in April 1996.
45. I understand that, as part of a litigation that has since settled, Kevin and I were accused of having both publicly used our technology and offered it for sale more than one year before the filing date of the '061 Patent; that is, before October 29, 1995. I have no recollection of any targeted ad serving by us prior to that date, and, as discussed above, the documents I have reviewed confirm for me that we did not have the ability to serve targeted ads before December 1995.
46. We did in fact test targeting of ads in December 1995. These ads were served for Internet Security Systems and IAN. These ads were served as part of a test to see if we could, in fact, serve targeted ads over the Internet. IAN received no money from ISS for serving these ads and these ads did not represent a commercial sale. (See Exhibit G.)
47. If we had used any aspect of our technology on the Internet prior to October 29, 1995, it would have been purely in an experimental way. Even in January, we were still testing the system. We had only developed the basic ideas of the invention in September – as discussed above, an immense amount of coding and development took place over the next few months.
48. For similar reasons, I do not believe we would have – or could have – offered anything we were developing for sale prior to October 29, 1995. Our technology was still developing, still experimental, and still far away from a production system. We were still working out the details of how the system would work. I firmly believe that we did not have a product or service that was ready for sale on that date.
49. I further declare that all statements made herein of my knowledge are true and that all statements made on information and belief are believed to be true; and further that the undersigned acknowledges that any false statements and the like so made are punishable by fine or imprisonment or both under §1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of any patent that issues from U.S. Reissue Application No. 09/577,798.

Date 2/9/01

  
Dwight A. Merriman